

Deformation Quantization of Distributed Inference: The Convex Case

Ryan J. Buchanan*
Parker M.D. Emmerson†

March, 2026

Abstract

Using modest spectral graph theory, we show that under the assumption of convexity, beliefs will diffuse towards consensus. Our toy model captures opinion dynamics in a manner sensitive to the order of belief-updates amongst agents in a network. To accomplish this, we introduce a first-order deformation of the classical observable algebra and study the resulting non-commutative correction through an explicit graph bracket. We include a concrete computation alongside some code in the appendix.

Contents

1	Introduction	1
2	Gradient Flow of Belief Configurations	4
3	Explicit Deformation: A Poisson-Type Graph Star Product	5
4	The Classical Limit of Consensus	7
A	Computational Illustration	8

1 Introduction

Graphs are the cornerstone of modern social network analysis. While graph clustering [11] has been used to identify cliques in social networks, allowing us to numerically probe real datasets, more abstract formulations of agentive dynamics have been put forth using factor graphs [7] or hypergraphs [12]. The simplicity of these structures allows one to formulate networks with minimal effort.

To illustrate this, let $\mathcal{N}$ be a network of sensors, agents, or distributed processors. We suppose that $\mathcal{N}$ has the structure of a graph

$$\mathcal{N} := (V, E), \tag{1.1}$$

*Active Inference Institute, Cascadia, USA 0009-0000-1253-0031

†Independent Researcher 0009-0007-1288-3292

where V is the set of vertices and E is the set of communication links. Throughout, we assume that communication is local in the sense that a vertex v_i may directly communicate with a vertex v_j only when there exists an edge $e_{ij} \in E$ joining them. This gives rise to an adjacency relation

$$v_i \sim v_j. \quad (1.2)$$

Although one may distinguish between noisy and noiseless channels in distributed inference [8], we will not do so here.

The graph Laplacian provides a natural quantitative proxy for disagreement across the network.¹ Let D denote the degree matrix of $\mathcal{N}$ and A its adjacency matrix. The Laplacian is then

$$L = D - A. \quad (1.3)$$

If $f : V \rightarrow \mathbb{R}$ is a real-valued function on the vertices, then

$$(Lf)(i) = \sum_{j \sim i} (f(i) - f(j)). \quad (1.4)$$

The associated quadratic form is

$$Q(f) := f^T L f = \sum_{(i,j) \in E} (f(i) - f(j))^2. \quad (1.5)$$

We interpret the function f as a local assignment of internal states to the agents of $\mathcal{N}$. For simplicity, these internal states will be called *beliefs*. In this interpretation, the Laplacian measures the extent to which a given belief configuration fails to agree with its neighboring values. Accordingly, the quadratic form $Q(f)$ measures the total disagreement across the network.

This suggests a natural energetic interpretation. If all local beliefs were perfectly harmonized, then the disagreement functional would vanish. Hence $Q(f)$ may be viewed as the amount of energy, effort, or dynamical correction required in order to transport the present configuration to a state of global consistency. In particular, when the vertices of $\mathcal{N}$ are interpreted as measuring devices, one may regard $Q(f)$ as the deviation of observed values from an idealized classical profile.

Definition 1.1. *Let*

$$\mathcal{A}_0(\mathcal{N}) := \{f : V \rightarrow \mathbb{R}\} \quad (1.6)$$

denote the classical algebra of local belief observables on $\mathcal{N}$.

The algebra $\mathcal{A}_0(\mathcal{N})$ is commutative under pointwise multiplication. Explicitly, if $f, g \in \mathcal{A}_0(\mathcal{N})$, then

$$(fg)(i) := f(i)g(i). \quad (1.7)$$

We interpret this product as the naive coexistence of two local inferences at the same vertex. In other words, the classical product does not register any contextual interference, order-sensitivity, or transport effect between distinct belief updates.

Remark 1.1. *The commutativity of $\mathcal{A}_0(\mathcal{N})$ makes it too rigid for modeling genuinely contextual inference. If one first incorporates the information carried by f and then that carried by g , the classical product does not distinguish this from the reverse order. A deformation of the commutative algebra $\mathcal{A}_0(\mathcal{N})$ is therefore a natural next step.*

¹For the standard reference, see [3].

Following the general philosophy of deformation quantization [4, 10], we introduce a formal parameter $\hbar$ and consider the formal algebra

$$\mathcal{A}_\hbar(\mathcal{N}) := \mathcal{A}_0(\mathcal{N})[[\hbar]]. \quad (1.8)$$

The role of $\hbar$ here is not necessarily that of Planck's constant in the physical sense. Rather, it will serve as a bookkeeping parameter which measures the strength of nonclassical correction to naive belief-combination.

Definition 1.2. A deformation coupling on $\mathcal{N}$ is a star product

$$\star : \mathcal{A}_\hbar(\mathcal{N}) \times \mathcal{A}_\hbar(\mathcal{N}) \longrightarrow \mathcal{A}_\hbar(\mathcal{N}) \quad (1.9)$$

of the form

$$(f \star g)(i) = f(i)g(i) + \hbar B_1(f, g)(i) + \hbar^2 B_2(f, g)(i) + \cdots, \quad (1.10)$$

where the B_k are bidifferential operators.

The distinction between $Q(f)$ and $\hbar$ is important. The functional $Q(f)$ measures disagreement in a given network configuration, whereas $\hbar$ controls the formal deformation of the algebra of observables. One may think of the geometry induced by the graph Laplacian as supplying the background structure from which the higher correction terms B_k are built, while $\hbar$ records the order at which those corrections occur.

Remark 1.2. In the inference-theoretic reading, the star product is interpreted as a deformed rule for composing local beliefs. The zeroth-order term $f(i)g(i)$ describes naive joint inference, while the higher-order corrections $B_k(f, g)(i)$ encode contextual interaction, transport effects, or hierarchy-sensitive couplings induced by the communication structure of the network.

This point of view also allows one to distinguish between static coexistence and dynamical correction. The classical term $f(i)g(i)$ represents the simultaneous presence of two local observables, whereas the operators B_k measure the way in which these observables influence one another once communication constraints are taken into account. In a loose but suggestive sense, the terms B_k may therefore be interpreted as generalized kinetic corrections to classical inference.

If the communication network is perfectly coherent, then the disagreement energy $Q(f)$ should vanish. In this limit, one expects the classical picture to be recovered. As a heuristic ansatz, we may therefore model an effective measurement by

$$M(f) = O^{\text{cl}} + Q(f), \quad (1.11)$$

where O^{cl} denotes the classical observable and $Q(f)$ measures the correction induced by disagreement across the network. This does not yet constitute a full quantization scheme, but it motivates the idea that disagreement energy should deform the effective composition law of observables.

The noncommutativity of the star product is of particular interest in the inference picture. Since, in general,

$$f \star g \neq g \star f, \quad (1.12)$$

the order in which local beliefs are updated begins to matter. This leads to the following quantity.

Definition 1.3. The star commutator of two observables $f, g \in \mathcal{A}_\hbar(\mathcal{N})$ is defined by

$$[f, g]_\star := f \star g - g \star f. \quad (1.13)$$

Proposition 1.1. *If the star product on $\mathcal{A}_\hbar(\mathcal{N})$ is noncommutative, then local belief-updates on $\mathcal{N}$ are order-sensitive.*

Proof. This is immediate from the definition. If $f \star g \neq g \star f$, then composing the update associated to f with the update associated to g yields a different result from composing them in the reverse order. Hence the inference process depends on order. $\square$

Accordingly, the star commutator measures the failure of contextual inference to commute. In the present framework, it may be interpreted as a quantitative measure of *contextual inference incompatibility*: the larger $[f, g]_\star$ is, the further the system is from an order-independent regime of information fusion.

The guiding idea of this note is thus simple. A communication network carries a classical disagreement geometry through its graph Laplacian, and this geometry may be used to motivate a deformation of the commutative algebra of local beliefs. The resulting star product provides a noncommutative rule for combining local inferences, while the corresponding star commutator detects the extent to which such combinations depend upon order, context, and communication structure.

Remark 1.3. *The present construction should be read as a first-order toy model of deformation-based distributed inference on graphs, intended to isolate the emergence of order-sensitive belief composition in the convex consensus regime rather than to provide a complete deformation-quantization theorem for graph observables.*

2 Gradient Flow of Belief Configurations

The variational formulation of disagreement naturally induces a dynamics on the space of belief configurations.

Definition 2.1. *The Inference Action associated to a belief configuration $f \in \mathcal{A}_0(\mathcal{N})$ is defined by*

$$S(f) := \frac{1}{2} \langle f, Lf \rangle = \frac{1}{2} Q(f). \quad (2.1)$$

Definition 2.2. *Let $S(f) = \frac{1}{2} \langle f, Lf \rangle$ be the inference action. The gradient flow of beliefs is defined by the evolution equation*

$$\frac{df}{dt} = -\nabla S(f). \quad (2.2)$$

With respect to the inner product $\langle \cdot, \cdot \rangle$, the gradient of S is given by

$$\nabla S(f) = Lf, \quad (2.3)$$

so that the flow takes the explicit form

$$\frac{df}{dt} = -Lf. \quad (2.4)$$

Property 2.1. *Along the gradient flow, the inference action is non-increasing:*

$$\frac{d}{dt} S(f(t)) = -\langle Lf, Lf \rangle \leq 0. \quad (2.5)$$

Proof. Differentiating $S(f(t))$ along the flow gives

$$\frac{d}{dt}S(f(t)) = \langle \nabla S(f), \frac{df}{dt} \rangle = \langle Lf, -Lf \rangle = -\langle Lf, Lf \rangle \leq 0, \quad (2.6)$$

using the positive-definiteness of the inner product. $\square$

Corollary 1. *The flow converges toward configurations satisfying*

$$Lf = 0, \quad (2.7)$$

*i.e., harmonic belief configurations.*²

Remark 2.1. *The functional $S(f)$ may be interpreted as a free-energy-like quantity [9, 5] for the network. In this interpretation, the gradient flow*

$$\frac{df}{dt} = -Lf$$

corresponds to a relaxation dynamics in which belief configurations evolve toward lower-energy states. We emphasize, however, that $S(f)$ captures only a structural notion of disagreement, and does not incorporate entropic or probabilistic contributions present in full variational free energy formulations.

These equilibrium states correspond to configurations of minimal disagreement across the network. In particular, when the underlying graph is connected, the kernel of L consists of constant functions, so that the flow drives the system toward consensus.




Figure 1: Gradient flow of belief configurations under $\dot{f} = -Lf$. Initial disagreement dissipates over time, leading to a uniform consensus state in which the deformation becomes effectively commutative.

Remark 2.2. *The evolution equation $\frac{df}{dt} = -Lf$ is formally analogous to a diffusion or heat equation on the network. In this interpretation, disagreement propagates and dissipates over time, leading to progressively smoother belief configurations.*

Remark 2.3. *This provides a concrete realization of the variational principle suggested in the previous section: inference dynamics may be modeled as a gradient descent process minimizing the disagreement action. In particular, the Laplacian acts as the generator of belief equilibration across the network.*

²For the introduction of *epistemic harmony*, see our earlier paper, [1]. For an unrelated notion, see [2].

3 Explicit Deformation: A Poisson-Type Graph Star Product

To ground the formal algebra of Section 1, we provide an explicit construction of the first-order correction $B_1(f, g)$. In the present note, this correction is taken to be a global scalar bilinear bracket, rather than a vertexwise observable. A natural candidate is a Poisson-type bracket capturing how beliefs interact when propagated across oriented edges.

Definition 3.1. Let $d : \mathcal{A}_0(\mathcal{N}) \rightarrow \mathbb{R}^{|E|}$ be the graph coboundary operator, where $(df)_{ij} = f(j) - f(i)$. Given an orientation or "flow preference" $\omega \in \mathbb{R}^{|E|}$ on the network, we define the first-order bilinear bracket

$$B_1(f, g) := \{f, g\}_\omega = \sum_{(i,j) \in E} \omega_{ij} ((df)_{ij} g(i) - (dg)_{ij} f(i)). \quad (3.1)$$

This produces a scalar first-order correction determined by the oriented edge geometry of the network.

Remark 3.1. The bracket $\{f, g\}_\omega$ is bilinear and antisymmetric by construction. We do not claim that it satisfies the Jacobi identity in full generality, but only that it serves as a first-order Poisson-type correction governing the deformation at the level of a global interaction functional. The weights ω_{ij} encode a preferred orientation or flow structure on the network, thereby introducing an intrinsic asymmetry into the interaction.

Example 3.1. Consider the 3-cycle network with oriented edges

$$(v_1, v_2), \quad (v_2, v_3), \quad (v_3, v_1),$$

and take $\omega_{12} = \omega_{23} = \omega_{31} = 1$. Define

$$f(v_1) = 1, \quad f(v_2) = f(v_3) = 0, \quad g(v_2) = 1, \quad g(v_1) = g(v_3) = 0.$$

Thus f is localized at v_1 and g at v_2 .

Classically, the pointwise product vanishes:

$$(fg)(v_i) = 0 \quad \text{for all } i,$$

since the supports of f and g are disjoint.

For the deformed term, we compute

$$B_1(f, g) = \sum_{(i,j) \in E} ((df)_{ij} g(i) - (dg)_{ij} f(i)).$$

Now

$$(df)_{12} = f(v_2) - f(v_1) = -1, \quad (dg)_{12} = g(v_2) - g(v_1) = 1,$$

while

$$(df)_{23} = 0, \quad (dg)_{23} = -1, \quad (df)_{31} = 1, \quad (dg)_{31} = 0.$$

Hence the three edge contributions are

$$(v_1, v_2) : (-1) \cdot g(v_1) - (1) \cdot f(v_1) = -1,$$

$$(v_2, v_3) : 0 \cdot g(v_2) - (-1) \cdot f(v_2) = 0,$$

$$(v_3, v_1) : 1 \cdot g(v_3) - 0 \cdot f(v_3) = 0.$$

Therefore

$$B_1(f, g) = -1.$$

So, to first order in $\hbar$,

$$f \star g = fg + \hbar B_1(f, g) = -\hbar.$$

By antisymmetry, $B_1(g, f) = 1$, so

$$g \star f = \hbar.$$

This exhibits explicitly that the deformed product distinguishes the order of belief composition, even though the classical product vanishes.

Remark 3.2. In this setting, $\hbar$ acts as an effective communication range. Even if two agents do not share a vertex, the deformed algebra allows their inferences to interact through higher-order corrections, modeled as a transport effect induced by the underlying graph geometry.

4 The Classical Limit of Consensus

We now link the noncommutative algebra back to the gradient flow defined in Section 2. We show that as the network approaches equilibrium, the contextual interference measured by the star commutator vanishes.

Theorem 4.1 (Recovery of Commutativity). *Let $f(t)$ and $g(t)$ be configurations evolving under the gradient flow $\dot{\phi} = -L\phi$. For a connected graph $\mathcal{N}$, the star commutator vanishes in the infinite-time limit:*

$$\lim_{t \rightarrow \infty} [f(t), g(t)]_\star = 0. \quad (4.1)$$

Proof. From Corollary 1, we know that for a connected graph,

$$\lim_{t \rightarrow \infty} f(t) = c_f \mathbb{1}, \quad \lim_{t \rightarrow \infty} g(t) = c_g \mathbb{1},$$

for some constants $c_f, c_g \in \mathbb{R}$.

Since each bidifferential operator B_k depends only on local differences of its arguments (in particular, B_1 is built from edge differences), it vanishes when evaluated on constant functions. Thus,

$$\lim_{t \rightarrow \infty} (f \star g) = c_f c_g \mathbb{1} + \hbar(0) + \hbar^2(0) + \dots \quad (4.2)$$

$$= c_f c_g \mathbb{1}. \quad (4.3)$$

Because the classical product is commutative, the commutator $[f, g]_\star$ vanishes in the limit. $\square$

Remark 4.1. Consensus realizes a classical, commutative regime of the model. Order-sensitivity and contextual incompatibility are sustained only in the presence of nontrivial disagreement energy $Q(f)$. As the network equilibrates, the higher-order deformation terms vanish, and the algebra reduces to its commutative limit.

Remark 4.2. The convergence to consensus described above reflects the convex, diffusive nature of the present model, in which disagreement is uniformly penalized. In more general settings—such as weighted, nonlinear, or constrained interactions—one may instead obtain locally coherent or multi-cluster equilibria. The present results should therefore be understood as describing the convex baseline regime of distributed inference.

Future Work

The present formulation models belief configurations as scalar-valued functions on the vertex set of a network. This corresponds implicitly to a trivial bundle with fiber $\mathbb{R}$ over the graph. A natural extension is to replace this scalar structure with a fibered one.

In such a setting, each vertex $i \in V$ is equipped with a local space of beliefs $\mathcal{F}_i$, and a configuration becomes a section of a bundle over the network. Interactions between agents are then governed not by simple differences $f(j) - f(i)$, but by a notion of parallel transport along edges. Concretely, one introduces transport operators

$$T_{ij} : \mathcal{F}_i \rightarrow \mathcal{F}_j,$$

allowing for the comparison of beliefs in different local contexts.

The resulting generalized gradient takes the form

$$\nabla_{ij} f = f(j) - T_{ij} f(i),$$

leading to a covariant Laplacian and a corresponding extension of the disagreement functional.

In this framework, inconsistency across the network is naturally captured by curvature: transporting a belief around a closed loop need not return it to its original value. Such holonomy effects provide a geometric measure of contextual incompatibility, extending the scalar notion of disagreement energy.

From the perspective of deformation, one may expect the star product to interact nontrivially with the bundle structure, potentially leading to a deformation theory of sections of nontrivial bundles over graphs.

These directions suggest a broader geometric framework in which distributed inference is governed by connections, curvature, and transport, rather than purely scalar diffusion. We leave the systematic development of this theory to future work.

Finally, the work remains to be extended to the non-convex case. This analysis should yield some genuinely surprising results, and will broaden the applicability of this framework. Since there do not always exist geodesics between two points in a non-convex space, this means that energy minimization, and thus perfect diffusion, cannot be assumed.

A Computational Illustration

To illustrate the gradient flow dynamics, we simulate the discrete-time update

$$f_{t+1} = f_t - \eta L f_t$$

on a simple cycle graph. Starting from a non-uniform initial configuration, the system converges to a constant state, confirming the theoretical prediction that the Laplacian flow drives the network toward consensus.

The following Python code illustrates the convergence of beliefs in a simple two-agent setup:

```
import numpy as np
import networkx as nx
import matplotlib.pyplot as plt
import scipy
```



```

# Create a simple graph (triangle)
G = nx.cycle_graph(3)

# Laplacian matrix
L = nx.laplacian_matrix(G).toarray()

# Initial belief configuration
f = np.array([1.0, 0.0, 2.0])

# Time step
eta = 0.1
steps = 30

history = [f.copy()]

# Gradient flow iteration:  $f_{t+1} = f_t - \eta * L f_t$ 
for _ in range(steps):
    f = f - eta * L @ f
    history.append(f.copy())

history = np.array(history)

# Plot evolution
for i in range(3):
    plt.plot(history[:, i], label=f'Node {i}')

plt.title("Gradient Flow to Consensus")
plt.xlabel("Time step")
plt.ylabel("Belief value")
plt.legend()
plt.show()

```

References

- [1] R. J. Buchanan and P. M. D. Emmerson *An Application of Modal Logic to Costello-Gwilliam Factorization Algebras* (2025) DOI: 10.22541/au.175131320.01270086/v2
- [2] C. G. Butler *SpiralOS: VI - Harmonic Completion - The First Breath* (2025) Available: <https://philarchive.org/rec/BUTSVZ>
- [3] F. R. K. Chung. *Spectral Graph Theory*. American Mathematical Society. (1997)
- [4] K. Costello and O. Gwilliam. *Factorization algebras in quantum field theory*. (2016)
- [5] K. Friston. *Life as we know it*. Journal of the Royal Society Interface, 10(86):20130475, (2013). DOI:10.1098/rsif.2013.0475
- [6] J. Hansen and R. Ghrist. *Opinion Dynamics on Discourse Sheaves*. (2020). <https://doi.org/10.48550/arXiv.2005.12798>
- [7] C.-S. Hu. *A Sheaf-Theoretic and Topological Perspective on Complex Network Modeling and Attention Mechanisms in Graph Neural Models* (2026) arXiv:2601.21207
- [8] S. Kar, S. Aldosari, J. M. F. Moura. *Topology for Distributed Inference on Graphs*. (2006). arXiv:cs/0606052
- [9] V. B. Knight, R. J. Cordes, D. A. Friedman. *The Free Energy Principle & Active Inference: a Systematic Literature Analysis* (2022) 10.5281/zenodo.7449368
- [10] M. Kontsevich. *Deformation quantization of Poisson manifolds, I*. (1997). arXiv:q-alg/9709040
- [11] R. Liu, S. Feng, R. Shi, W. Guo. *Weighted Graph Clustering for Community Detection of Large Social Networks*. Procedia Computer Science Vol. 31 (2014) <https://doi.org/10.1016/j.procs.2014.05.248>
- [12] T. Mukande, E. Ali, A. Caputo, R. Dong, N. O'Connor *Towards Efficient Hypergraph and Multi-LLM Agent Recommender Systems* (2025) arXiv:2512.06590